



CONTENT DELIVERY NETWORK USING DIFFERENTIAL CACHING

BACKGROUND OF THE INVENTION

Field of the Invention

[001] This invention relates to a content delivery network using differential caching.

Related Art

[002] When multiple users (at client devices) request information from a server (at a server device), it often occurs that the number of requests from client devices taxes the server device, and reduces the quality of service that each user experiences. Moreover, when those multiple client devices are distributed at widely disparate locations, there is reduced quality of service experienced by users relatively distant from the server device, due to distance (either measured physically or measured by a communication network topology) that messages travel. Accordingly, it would be advantageous to provide additional server devices having the same content for delivery to client devices, to (1) share the load of requests made to the server device, and to (2) move the content for delivery closer to client devices. Each of these effects should improve the quality of service experienced by client devices.

[003] One known method is to provide a content delivery network, including an originating server device and a set of mirroring server devices, disposed so that original content from the originating server is delivered and maintained at each of the mirroring servers. While this known method generally achieves the goal of moving content for delivery closer to client devices, it has the drawback that it is relatively unsuitable for content that is not static. When the content for delivery is dynamically changing, or is personalized for users at different client devices, the content is not static, and the mirroring servers cannot guarantee that they have the correct content for delivery. The content delivery network thus is relatively unsuitable for responding to requests for non-static content.

[004] Accordingly, it would be desirable to provide a technique for serving relatively non-static content for delivery in a content delivery network.

SUMMARY OF THE INVENTION

[005] The invention provides a method and system for serving relatively non-static content for delivery in a content delivery network. Content for delivery is pushed out from an originating server to a set of mirroring servers using differential caching. Using differential caching, each object that might be requested by a client is separated into template information and delta information. In a preferred embodiment, the originating server determines a set of templates for differential caching, and causes those templates to be distributed using the content delivery network to the mirroring servers. Each mirroring server in the content delivery network is able to access, either locally or by reference to a nearby distribution point, a copy of an applicable recent template for each request made by a client. Hosting of the template information is decoupled from hosting of the delta information.

[006] Delta information can include (a) changes to the template information since last distributed using the content delivery network, or (b) personalized information specific to the client or to the request made by the client. In a preferred embodiment, delta information can be served separately from the originating server, or can be served separately from mirroring servers using a content distribution network (possibly the same content distribution network as used to serve template information).

[007] Using differential caching with the content distribution network (that is, decoupling distribution of template information from distribution of delta information) allows the system to provide the following functions, not provided by known methods:

- The amount of bandwidth used to distribute changes to content from the originating server is minimized. This allows more content to be distributed using the same amount of bandwidth, or alternatively, reduces the cost of distributing the same amount of content.
- The content distribution network is able to distribute dynamically changing content with substantially less use of bandwidth, and with substantially greater responsiveness to the dynamic changes. The template

information is sufficiently static for distribution using the content distribution network. Distribution of the delta information does not put an excessive burden on the originating server, or when the delta information is distributed using a content distribution network, does not put an excessive burden on the content distribution network.

- The content distribution network is able to distribute personalized content. As with dynamically changing content, the template information is sufficiently static for distribution using the content distribution network. Distribution of the delta information can be distributed from the same server or a different set of servers. For example, if personalized content requires a database lookup, the originating server or another server can perform primarily database lookups, while a separate mirroring system for the personalized information can distribute the delta information to clients in personalized form.
- There are many individual objects that might be included in a page that is specific to a particular client. While these individual objects are themselves typically very similar or even identical to objects to be included in pages specific to different clients, the selection of which individual objects to be included is responsive to personalized information about each client. Thus, distribution of the individual objects can be decoupled from personalization using those objects, just like distribution of the template information can be decoupled from distribution of delta information.

[008] The invention can be used in conjunction with a known content distribution network, so as to allow the known content distribution network to provide both dynamic content and personalized content (when used in conjunction with the invention).

[009] The invention has general applicability to content delivery, not limited specifically to the web pages, web protocols, or caching (and not limited specifically to content delivery as described herein). For example, embodiments of the

invention can include one or more of, or some combination of, the following applications:

- distribution of databases responses, including responses to common or frequently used database queries;
- distribution of email and groupware messages, bulletin board or newsgroup messages; and
- distribution of streaming media content (template information and insertion points can be streamed from a content delivery network and delta information can be streamed from an originated server)

[010] Moreover, techniques used by a preferred embodiment of the invention for content delivery can be used in contexts other than the specific applications disclosed herein. For example, techniques used by embodiments of the invention for content delivery are all generally applicable to fields other than the specific applications disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[011] Figure 1 shows a block diagram of a content delivery network using differential caching.

[012] Figures 2A and 2B show a process flow diagram of a method for operating a content delivery network using a proxy encoder server and one or more decoders.

[013] Figure 3 shows a process flow diagram for a method of operating a content delivery network using differential caching.

DESCRIPTION OF THE PREFERRED EMBODIMENT

[014] The invention is described herein with regard to preferred steps and data structures. Those skilled in the art will recognize, after perusal of this application, that the described steps and data structures are not limited to any particular processing devices (whether general-purpose or special-purpose processing devices, or specific circuitry). Rather, those of ordinary skill in the art would be able to implement the described steps and data structures, and equivalents

thereof, without undue experimentation or further invention. All such implementations are within the scope and spirit of the invention.

[015] Inventions described herein can be used in conjunction with inventions described in the following applications:

- Application Serial No 09/436,136, filed November 11, 1999, in the name of Stephane KASRIEL, titled "Predictive Pre-Download of Network Objects", issued as U.S. Patent No. 6,721,780 on April 13, 2004;
- Application Serial No. 09/734,910, filed December 11, 2000, in the name of Stephane KASRIEL, titled "Predictive Pre-Download Using Normalized Network Objects";
- Application Serial No. 09/827,268, filed April 4, 2001, in the name of Stephane KASRIEL, titled "Server-Originated Differential Caching";

[016] Each of these applications is hereby incorporated by reference as if fully set forth herein. They are collectively referred to as the "incorporated disclosures".

[017] Lexicography

- **client** and **server** – as used herein, the phrases, "**client**" and "**server**" refer to a relationship between two devices, particularly to their relationship as client and server, not necessarily to any particular physical devices.
- **client device** and **server device** – as used herein, the phrase "**client device**" includes any device taking on the role of a client in a client-server relationship (such as an HTTP web client and web server). There is no particular requirement that any client devices must be individual physical devices; they can each be a single device, a set of cooperating devices, a portion of a device, or some combination thereof. As used herein, the phrase "**server device**" includes any device taking on the role of a server in a client-server relationship. There is no particular requirement that server devices must be individual physical devices; they can each be a single device, a set of cooperating devices, a portion of a device, or some combination thereof.

- **logically remote** – as used herein, the phrase “**logically remote**” refers to the relative logical placement or degree of connectivity between two or more computerized systems or two or more elements within a single system. Generally, elements that are relatively proximate to each other may be logically remote if there is a small probability that information will flow between them on a regular basis.

[018] System Elements

[019] Figure 1 shows a block diagram of a content delivery network using differential caching.

[020] A system includes one or more clients 110, an originating server 120, a set of mirroring servers 130, a set of proxy encoder servers 140 and a communication network 150.

[021] Client

[022] Each client 110 includes a client workstation 111 and a client operator 112.

[023] As described in the incorporated disclosures, a “workstation” might include a personal computer, a software package on a server, a handheld computer cooperating with a personal computer or with a server (or both), or a telephone interface to a system such as an interactive voice response system. There is also no particular requirement that multiple workstations used by a single client need be of the same type. Also as described in the incorporated disclosures, an “operator” might comprise an individual person, a set of persons having authority to act in particular way, a proxy for an individual person or set of persons, such as a human secretary or a computer program having the function of forwarding or aggregating or scheduling requests made by others, or even an artificial intelligence program such as an expert system or otherwise.

[024] In a preferred embodiment, each client 110 includes a web browser 113, such as the “Internet Explorer” product or the “Netscape Navigator” product, capable of using a message transfer protocol, such as HTTP (hypertext transfer protocol), or a variant thereof, to request documents (such as for example web pages) from

the originating server 120 or a mirroring server 130 and to receive documents and other responses from the originating server 120 or a mirroring server 130. A decoder 114 is coupled to the web browser 113, preferably as a browser add-on. However, in other embodiments, multiple decoders 114 may be situated relatively proximate to an Internet service provider, an enterprise cache or at other locations within the communications network 150.

[025] In other embodiments, the browser 113 is not coupled to a decoder 114. In such embodiments, functions normally performed on the client side by the decoder 114 are performed by the browser 113 and the proxy encoder server 140. These embodiments are referred to as "clientless versions".

[026] Originating Server

[027] The originating server 120 includes a computer 121 and a database 122 of documents 123. In a preferred embodiment, documents 123 can include (as further described herein) web pages, embedded objects for web pages, template web pages, changed data for insertion into template web pages, and code fragments.

[028] The originating server 120 includes a processor, program and data memory, and operates under control of software to perform the tasks described herein. In particular, the originating server 120 is capable of using a message transfer protocol, such as HTTP or a variant thereof, to receive requests for documents (such as for example web pages) from clients 110 and to respond to those requests by sending those documents to clients 110. In a preferred embodiment, the originating server 120 uses HTTP version 1.1, or at least some features thereof, as described herein.

[029] Mirroring Servers

[030] Similar to the originating server 120, the mirroring servers 130 each include a computer 131 and a database 132 of documents 133. Similar to the originating server 120, each mirroring server 130 includes a processor, program and data

memory, and operates under control of software to perform the tasks described herein.

[031] Proxy encoder servers

[032] Similar to the originating server 120 and the mirroring servers 130, the proxy encoder servers 140 each include a computer 141 and a database 142 of documents 133 or template information 124. However, the proxy encoder server 140 also includes a software element 143 that recognizes a URL and transparently alters the URL so as to direct requests from the client 110 to the content delivery network. This software element 143 can also embed information in the URL such as may be relevant to which version of a web page or template information is transmitted, whether a decoder 114 can accept delta information 125 and other information that pertains to equipment or communication parameters.

[033] In a preferred embodiment, the proxy encoder server 140 is relatively local to the originating server 120 and the mirroring servers 130. The encoder 140 may also be integrated inside an existing server component, such as any of the previously described servers, a web server, an application server, a cache or an L7 switch.

[034] It would be clear to those skilled in the art, after perusal of this application, that the system can include more than one originating server 120, each disposed for originating a set of content for distribution (probably a completely different set of such content for distribution) to clients 110. Moreover, it would be clear to those skilled in the art, after perusal of this application, that at least some of the originating servers 120 can operate as mirroring servers 130 for other, different, originating servers 120, while at least some of the mirroring servers 130 can operate as originating servers 120 for their own specific content for distribution. All such variations and extensions of ideas presented herein are within the scope and spirit of the invention, and would not require either undue experimentation or invention when considered in view of this application.

[035] Communication Network

[036] Clients 110, originating servers 120 and mirroring servers 130 are coupled using a communication network 150. In a preferred embodiment, the communication network 150 includes a computer communication network, such as the Internet. However, in alternative embodiments, the communication network 150 might include an intranet, extranet, VPN (virtual private network), ATM system, a portion of a private or public PSTN (public switched telephone network), a frame relay system, or any other communication technique capable of performing the functions described herein.

[037] Content Distribution

[038] In a preferred embodiment, each mirroring server 130 performs a function much like a caching device. Each mirroring server 130 receives requests from clients 110, and determines whether (a) it can service those requests directly, or (b) it can service those requests by obtaining information from the originating server 120. When the mirroring server 130 can service those requests directly, it does so, sending content for delivery to the requesting client 110 in response to the request. When the mirroring server 130 can service those requests by obtaining information from the originating server 120, it does so, requesting the content from the originating server 120, and sending that content to the requesting client 110. The mirroring server 130 can determine whether to cache (or otherwise maintain) a copy of the content for delivery, so that later requesting clients 110 can be serviced without resort to a request to the originating server 120.

[039] When maintaining a document 123 (such as a web page), the originating server 120 determines a set of template information 124 and a set of delta information 125. In a preferred embodiment, the template information 124 indicates those portions of the document 123 that are relatively static, while the delta information 125 indicates those portions of the document 123 that are relatively dynamic.

[040] In a preferred embodiment, requests are made by the client 110 using the decoder 114. The decoder 114 forwards the request to the proxy encoder server

140, which in turn, transmits the request to the originating server 120. The original HTTP response is preferably altered so that additional information can be appended to it, such as whether template information 124 or an entire page was fetched, what version of the template information 124 was obtained and similar information pertaining to other communication parameters. This additional information can be used in subsequent requests, so as to determine the relative freshness of a page. The proxy encoder server 140 retrieves the document 123 and uses it to create an associated template 124. In this embodiment, the document 123 is integrated at either the client device 110 by the client 110 or at the proxy encoder server 140 (if the client device 110 does not accept delta information 125).

[041] In another preferred embodiment, a request from a client 110 can be associated with delta information 125. The delta information 125 is associated with a pointer to a template 124 that is available, either through the content delivery network (if the template is available there) or at the originating server 120.

[042] In an alternative embodiment, each mirroring server 130 maintains a copy of template information 124 for that particular document 123. When the document 123 is requested by a client 110, the mirroring server 130 can provide the template information 124 to the client 110 from its cache, while the client 110 obtains the delta information 125 from the originating server 120 (or from a content distribution network similarly disposed for distributing delta information 125). Sending the template information 124 from the originating server 120 to the mirroring servers 130 (and from the mirroring servers 130 to the decoder 114) is separate from sending the delta information 125 from the originating server 120 to the client 110.

[043] Method of Operation

[044] Figures 2A and 2B show a process flow diagram for a method of operating a content delivery network using a proxy encoder server and one or more decoders.

- [045] A method 200 includes a set of flow points and process steps as described herein.
- [046] Although by the nature of textual description, the flow points and process steps are described sequentially, there is no particular requirement that the flow points or process steps must be sequential. Rather, in various embodiments of the invention, the described flow points and process steps can be performed in a parallel or pipelined manner, either by one device performing multitasking or multithreading, or by a plurality of devices operating in a cooperative manner. Parallel and pipelined operations are known in the art of computer science.
- [047] At a flow point 210, a client 110 is ready to make a request for a document 123 from a mirroring server 130 or an originating server 120 (depending upon which is closest). In a preferred embodiment, each request for a document 123 is performed independently, even if a plurality of requests are to be performed substantially simultaneously.
- [048] At a step 211, the client 110 generates a request message 151 to the mirroring server 130 or the originating server 120 for the document 123. The request message 151 identifies the document 123 and requests that the mirroring server 130 or the originating server 120 send the document 123 to the client 110. In a preferred embodiment, the request message 151 is made using the decoder 114, preferably on the client's web browser 113.
- [049] In a step 212, the decoder 114 forwards this request message 151 to the proxy encoder 140, situated preferably near the mirroring server 130 or the originating server 120. In forwarding this request, the decoder 114 indicates that the sender, (that is the client 110) is compatible with systems that provide delta encoding.
- [050] In clientless versions (that is, those clients without a decoder 114) the request goes directly to the encoder 140.
- [051] In a step 213, the proxy encoder server 140 fetches document 123 or a template 124 that corresponds to document 123 from either the originating server 120 or the mirroring server 130, depending upon which is closest to the proxy encoder 140. Upon obtaining the content, the proxy encoder 140 updates the template

124 for the document 123. In the event that there is not a template 124 associated with document 123, the proxy encoder 140 generates a template 124 and caches it.

- [052] In clientless embodiments, the proxy encoder 140 responds directly with the delta information. In such embodiments, the delta is an HTML page that includes a reference to a template (for example, a Javascript) and a delta (for example, some Javascript instructions). In this embodiment, the Javascript instructions comprising the delta tell the browser how to transform the template into the correct HTML or XML document.
- [053] In a step 214, the decoder 114 requests the template 124 from either the proxy encoder server 140 or the mirroring server 130, (depending where the encoder proxy 140 specified the template 124 was cached in the previous step, either immediately or at a later point in time, depending upon preferences set by the client operator 112).
- [054] In the clientless version, the browser 113 automatically and immediately retrieves the template from the site specified in the previous step.
- [055] In a step 215, the encoder proxy 140 sends the contents of the template 124, and a tag (also known as an "Etag") that corresponds to the version of the template. The encoder marks the template as being cacheable by network elements such as the mirroring servers 130, the client 110, the decoder 114 or public caches such as HTTP proxy caches.
- [056] If the contents and etag are sent from the mirroring server 130, then the mirroring server 130 searches its cache for the template 124. If the template 124 is present in its cache, the mirroring server 130 sends the template 124 directly to the client 110. If the template 124 is not present in the cache, the mirroring server 130 automatically fetches the template 124 from the encoder proxy 140, caches the fetched template 124, and sends the template to the client 110.
- [057] The following steps occur when the user subsequently requests the document 123.
- [058] In a step 216, the same or a different client 110 requests a document 123 by generating a request message 151. Similar to step 210 and 211, this is mediated

by the decoder 114 so that the request is directed to the proxy encoder 140. Continuing with the example, the decoder 114 changes the original request to as to further specify a version number that is used to ascertain if changes have occurred.

- [059] In a step 217, the proxy encoder 140 receives the request from the decoder 114. If the proxy encoder 140 determines that there is not a version of the template 124 in the database 142, the proxy encoder 140 obtains the document 123 from either the originating server 120 or the mirroring server 130 as in steps 213 – 215 and caches the template. If, however, the template 124 is available, the proxy encoder 140 calculates the differences between the version of template 124 that is available in the database 142 and a newer version of the document 123 such as may be available from the originating server 120 or the mirroring server 130. These differences are the delta information 125.
- [060] In a step 218, the proxy encoder 140 may send either the delta information 125 to the decoder 114 (that is, if the decoder 114 can accept delta information 125) or it sends the document 123 to the client.
- [061] Steps 216 through 218 are described with respect to a client – server implementation. In the clientless version, the proxy encoder server 140 does not need to know which version of the template 124 is at the client 110. This information is not needed because the proxy encoder server 140 makes this decision a priori and instructs the client 110 to use a specific version of the template 124. Under these circumstances, steps 216 – 218 in the clientless version are comparable to step 210 through 215 in which the browser 113 (rather than the decoder 114) makes a request that is directed to the encoder proxy 140.
- [062] Figure 3 shows a process flow diagram for a method of operating a content delivery network using differential caching.
- [063] A method 300 includes a set of flow points and process steps as described herein.
- [064] Similar to method 200, the flow points and process steps are described sequentially, there is no particular requirement that the flow points or process steps must be sequential. Rather, in various embodiments of the invention, the

described flow points and process steps can be performed in a parallel or pipelined manner, either by one device performing multitasking or multithreading, or by a plurality of devices operating in a cooperative manner. Parallel and pipelined operations are known in the art of computer science.

- [065] At a flow point 310, a client 110 is ready to make a request for a document 123 from a mirroring server 130. In a preferred embodiment, each request for a document 123 is performed independently, even if a plurality of requests are to be performed substantially simultaneously.
- [066] At a step 311, the client 110 generates a request message 151 (shown in figure 1) to the mirroring server 130 for the document 123. The request message 151 identifies the document 123 and requests that the mirroring server 130 send the document 123 to the client 110.
- [067] At a step 312, the mirroring server 130 determines if it has template information 124 for the requested document 123. To perform this step, the mirroring server 130 performs one the following sub-steps:
- At a sub-step 312(a), if the mirroring server 130 does not have the template information 124, it generates a request message 152 (shown in figure 1) to the originating server 120 for the template information 124. As part of this sub-step, upon receiving the template information 124, the mirroring server 130 records the template information 124 in a cache relatively local to the mirroring server 130.
 - At a sub-step 312(b), the mirroring server 130 has the template information 124, the method proceeds with the next step.
- [068] At a step 313, the mirroring server 130 generates a response message 153 (shown in figure 1) to the client 110 with the template information 124.
- [069] At a step 314, the client 110 (or the mirroring server 130, if the system is so configured) generates a delta request message 154 to the originating server 120 for the delta information 125. As part of this step, the system performs one of the following sub-steps:

- At a sub-step 314(a), if the client 110 generates the delta request message 154, the client 110 receives the delta information message 155 from the originating server 120
- At a sub-step 314(b), if the mirroring server 130 generates the delta request message 154, the mirroring server 130 receives the delta information message 155 from the originating server 120. As part of this sub-step, the mirroring server 130 forwards the delta information 125 in a second delta information message 155 to the client 110.

[070] At a step 315, the client 110 integrates the template information 124 and the delta information 125 into a unified document 123 for presentation. Integration of the template information 124 and the delta information 125 is further described in the incorporated disclosures.

[071] Generality of the Invention

[072] The invention has general applicability to content delivery, not limited specifically to the web pages, web protocols, or caching (and not limited specifically to content delivery as described herein). For example, embodiments of the invention can include one or more of, or some combination of, the following applications:

- distribution of databases responses, including responses to common or frequently used database queries; and
- distribution of email and groupware messages, bulletin board or newsgroup messages.

[073] Moreover, techniques used by a preferred embodiment of the invention for content delivery can be used in contexts other than the specific applications disclosed herein. For example, techniques used by embodiments of the invention for content delivery are all generally applicable to fields other than the specific applications disclosed herein.

[074] Other and further applications of the invention in its most general form would be clear to those skilled in the art after perusal of this application. The invention

would be usable for such other and further applications without undue experimentation or further invention.

[075] Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope and spirit of the invention; these variations would be clear to those skilled in the art after perusal of this application.